

Description

Method for managing a group of network access servers

5 The invention relates to a method for managing a group of network access servers, within which group the "Multichassis Multilink Point-to-Point Protocol" (MMP) is used, wherein an address list of the other network access servers in this group is managed by each network access server in this group. The
10 invention also relates to a network access server for carrying out the method according to the invention.

According to the prior art, individual connections in a packet data network using the Point-to-Point Protocol (PPP) can be
15 bundled into a virtual connection having increased bandwidth, wherein this virtual connection is called a Multilink PPP connection (MP) or MP bundle. The PPP has already been standardized in Network Working Group, Request for Comments: 1661, Category: Standards Track "The Point-to-Point Protocol
20 (PPP)", and the MP in Network Working Group, Request for Comments: 1990, Category: Standards Track "The PPP Multilink Protocol (MP)".

A typical application for this is e.g. the combination of two
25 B-channels of a basic-rate ISDN connection interface. The PPP packets belonging to the bundled connection can then be transmitted in parallel on both B-channels. Moreover, larger packets can be broken down and the resulting fragments distributed on the B-channels. In this case, a receiving MP
30 entity must be able to reassemble the received fragments and forward the packets in the correct sequence.

Since Remote Access Service (RAS) systems typically consist of a plurality of separate access modules, which are also called

Network Access Servers (NAS), a problem arises in the context of Multilink PPP. Specifically, the individual PPP connections belonging to an MP bundle are not generally routed to the same network access server. The information indicating whether an 5 individual PPP connection belongs to an MP bundle, and if so to which, is only available after completion of the PPP negotiations (the LCP phase) and the authentication phase.

It is therefore necessary for the individual network access 10 servers to notify each other concerning the MP bundles they are managing, this being known as "bundle discovery". Solutions already exist for this purpose, e.g. Network Working Group RFC 2701, Category: Informational Nortel Networks, "Multi-link Multi-node PPP Bundle Discovery Protocol" or Cisco, 15 "Multichassis Multi-link PPP (MMP)",
<http://www.cisco.com/warp/public/131/3.html>.

As part of the bundle discovery phase, one of the network access servers owning individual PPP connections for a specific 20 MP bundle is selected in order to assume the function of the so-called "bundle head". All other network access servers forward the multilink packets, which they receive on their PPP connections belonging to the MP bundle, to the bundle head. At the bundle head, the MP fragments are assembled and forwarded 25 in the correct sequence to the higher layers, e.g. to the Internet protocol.

The transport of the multilink packets between the network access servers takes place by means of a Layer-2 tunneling 30 protocol such as L2TP or L2F, for example. The described method is also called Multichassis Multilink PPP (MMP).

The group of network access servers within which MMP is required is also called a "stack group". The exact procedure

and the efficiency of the bundle discovery phase depend inter alia on whether the individual members know each other. This is not the case in the above-cited solution which is described in RFC 2701. If it is detected on a network access server that

5 Multilink PPP was negotiated on a local PPP connection, it is necessary to determine which of the following alternatives applies:

1. The MP bundle already exists and the local network access server itself is bundle head.
- 10 2. The MP bundle already exists and a different network access server is bundle head.
3. The MP bundle does not yet exist.

15 Whether case 1 applies can be detected locally on a network access server. In order to distinguish between case 2 and 3, however, it is necessary to perform a bundle discovery. The following procedure is described in RFC 2701 for this purpose: The network access server sends a request to a distribution 20 list address in the form of an IP Multicast. Since a network access server does not know which other members are included in the group, it is necessary to wait for a specific period in order to see whether a positive reply (in which case a bundle 25 is present) arrives from another network access server (the bundle head). Depending on the network topology concerned, the length of this period must be variously specified and can extend the time which is required to establish a connection.

If a network access server knows the other members of a group, 30 however, it is only necessary to wait until answers, both positive and negative, are received from all members. This is solved thus in the so-called Stack Group Bidding Protocol, (SGBP) (see also Cisco, "Multichassis Multilink PPP (MMP)",

<http://www.cisco.com/warp/public/131/3.html>), albeit with the disadvantage that the group must be manually configured.

5 The invention therefore addresses the problem of specifying an improved method for managing a group of network access servers.

According to the invention, this is achieved by means of a method of the type cited at the beginning, wherein the logging on/off of a network access server to or from this group takes 10 place in such a way that an address list of a network access server always indicates the current status of the network access servers in the group.

15 Since an address list is always updated when a network access server logs on/off, a configuration of the group of network access servers is essentially simplified for the operator of a packet data network. Various methods can be applied in this case, and these are cited in the dependent claims.

20 It is particularly advantageous if,

- when a new network access server logs onto a group of network access servers, a first message is sent from the new network access server to the network access servers of this group,
- the network access servers of this group store the address of the new network access server in an address list and send a second message to the new network access server in each case,
- the second messages are received and used by the new network access server for creating and storing an address list of all network access servers in this group.

30 This variant of the claimed method is particularly easy to implement and therefore particularly advantageous. In this

case, a first message of the new network access server is followed by second messages of the network access servers of a group. Since the messages contain the addresses of the senders, it is advantageously possible to establish address lists in the 5 network access servers, both in the new network access server and in the network access servers of the group.

It is also advantageous if

- a repetition time is assigned to a network access server in the group, said repetition time specifying the time intervals at which a second message is sent from the network access server in a periodically recurring manner to the other network access servers in the group, and
- the network access server is deleted from the address lists of the other network access servers in this group if the second message is not received by them before the expiry of the repetition time.

As part of this activity, a check advantageously determines 20 whether a network access server is actually still a member of the group of network access servers, or whether a connection to this server has failed due to a technical problem, for example. If this is the case, the relevant network access server is deleted from the address lists of the other network access 25 servers.

It is beneficial to provide a method in which

- the repetition time is contained in the first message, and
- this repetition time is stored in a list by the network 30 access servers of this group when a new network access server logs on.

In this variant, the repetition time is therefore directly transferred from the new network access server to the group of

network access servers when it logs on. The network access servers save the repetition time in succession in a list, and can advantageously begin to monitor the arrival of a second message immediately. In this case, it is conceivable to use a 5 separate list or a dedicated column in the address list.

It is also beneficial if, instead of the second message, a fourth message is provided for the periodically recurring notification. In this case, a second message is still used for 10 the logon procedure, but a fourth message which is independent of the logon procedure is used in order to check whether a network access server is actually still a member of the group of network access servers or whether e.g. a connection to this server has failed due to a technical problem. This is 15 advantageous in order to separate the individual method sections more effectively in relation also to the messages.

An advantageous variant of the invention also includes a method

- in which a third message is sent by one network access 20 server in the group to the other network access servers in the group and
- in which the other network access servers in this group delete this network access server from their address lists when they receive this message.

25

In this way, a network access server can be logged off from the group actively, i.e. not merely as a result of the absence of a second message. This advantageously avoids possible misunderstandings as to whether the absence of a second message 30 is due to an error or is intentional for the purpose of logging a network access server off.

It is advantageous if a distribution list address is used for sending the first and/or second and/or third messages and the

fourth messages within the group of network access servers, said distribution list address including addresses of at least all network access servers in this group, wherein a message contains an identification of the group.

5

The distribution list address therefore includes at least all potential members of a group. It is also conceivable for different groups to use the same distribution list address, and therefore the recipients of a message analyze a group 10 identification which is transferred at the same time. In this case, it is advantageous for the operator of the packet data network that only limited resources are required for the configuration of the network.

15 The problem addressed by the invention is also solved by means of a network access server which

- includes means for linking into a group of network access servers, wherein the "Multichassis Multilink Point-to-Point Protocol" (MMP) is used within said group,
- includes an address list of the other network access servers in this group, and
- includes means for registering a logging on and/or off of another network access server to or from this group, such that an address list in the network access server always 20 shows the current status of the network access servers which are in the group.

25

Since an address list is continuously updated when a network access server logs on and/or off, configuring the group of 30 network access servers is greatly simplified for the operator of a packet data network. It should be noted at this point that the advantages which are cited in relation to the claim method apply equally to the claimed network access server and vice versa.

In this case, it is particularly advantageous if the network access server includes

- means for receiving a first message, which message indicates the logging on of a new network access server to a group of network access servers,
- means for storing an address of the new network access server in an address list, wherein the address is contained in the first message,
- 10 - means for sending a second message to the new network access server, and
- means for receiving second messages and means for generating and storing an address list of all network access servers in a group, wherein the addresses are contained in the second messages.

This variant of the invention is particularly easy to implement and therefore particularly advantageous. In this case, a network access server comprises means for transferring the local address to the network access servers in a group, and means for establishing address lists of the other network access servers in a group.

It is advantageous to provide a network access server which includes

- means for the periodically recurrent sending of a second message to the other network access servers in the group,
- means for storing a repetition time which is assigned to a network access server,
- 30 - means for monitoring whether a second message of a network access server was received before the expiry of the repetition time which was assigned to it, and
- means for deleting a network access server from an address list.

In this context, the network access server advantageously includes elements for checking whether a network access server is actually still a member of the group of network access

5 servers or whether e.g. a connection to this server has failed due to a technical problem. Said network access server also includes means for continuously indicating its active participation to the other network access servers.

10 It is beneficial

- if the network access server includes means for storing a repetition time, which repetition time is contained in the first message, in a list.

15 A network access server according to the invention therefore includes means for directly storing the repetition time in a list when a new network access server logs on, and therefore it is advantageously possible to start to monitor the arrival of a second message immediately. In this case, it is conceivable to
20 use a separate list or a dedicated column in the address list.

It is also beneficial if

- the network access server includes means for sending a third message to the other network access servers in the

25 group, and

- the network access server includes means for deleting a network access server, from which server a third message has been received, from an address list.

30 In this case, a network access server advantageously includes means for detecting the active logging off of another network access server and means for indicating its own logging off to other network access servers. In this way, it is possible advantageously to avoid possible misunderstandings as to

whether the absence of a second message is due to an error or is intentional for the purpose of logging a network access server off.

5 A beneficial configuration of a network access server is also established if a distribution list address is provided as a means for sending the first and/or second and/or third and fourth messages, wherein said distribution list address includes addresses of at least all network access servers in
10 this group and contains an identification which has been assigned to the group.

The distribution list address therefore includes at least all potential members of a group. It is also conceivable for
15 different groups to use the same distribution list address, and therefore the recipients of a message have means for analyzing a group identification which is transferred at the same time. In this case, it is advantageous for the operator of the packet data network that only limited resources are required for the
20 configuration of the network.

The invention is explained in greater detail with reference to an exemplary embodiment which is illustrated in the figures and relates to the logging on and off of a network access server to
25 and from a group of network access servers, wherein:

Figure 1: shows the logging on of a new network access server ZSN+1 to a group of network access servers ZS1, ZS2, ..., ZSN;

30 Figure 2: shows an acknowledgement in the form of a second message N2 from the group of network access servers ZS1, ZS2, ..., ZSN to the new network access server ZSN+1;

Figure 3: shows the periodically recurring report from a network access server ZS to the other network access servers in the group;

Figure 4: shows the logging off of the n-th network access

5 server ZSN from the group of network access servers ZS.

Each of the Figures 1 to 4 shows a first to n-th access server ZS1, ZS2, ..., ZSN and a new access server ZSN+1. In this case,

10 all access servers ZS are connected to each other and can exchange messages over these connections. Furthermore, each access server ZS has a list of the other network access servers ZS in the group.

15 The function of the arrangement which is illustrated in the figures is explained in greater detail below.

Figure 1 shows the logging on of a new network access server ZSN+1 to a group of network access servers ZS1, ZS2, ..., ZSN.

20 Since each access server ZS includes a list of the other network access servers ZS in the group, the address list of the first network access server ZS1 shows e.g. the entries ZS2, ZS3, ..., ZSN. The address list of the new network access server ZSN+1 is still empty, however. As a first step, a first

25 message N1 is sent from the new network access server ZSN+1 to the group of network access servers ZS1, ZS2, ..., ZSN. This takes place via a distribution list address which includes the addresses of all network access servers in this group ZS1, ZS2, ..., ZSN. Since the distribution list address can also include

30 other network access servers ZS, the first message N1 contains an identification of the group GI. In addition, the first message N1 contains a repetition time TWN+1 which has been assigned to the new network access server ZSN+1.

Next - as shown in Figure 2 - the address of the new network access server ZSN+1 is entered in the address lists of the network access servers of the group ZS1, ZS2, ..., ZSN.

Therefore the address list of the first network access server

5 ZS1 now shows the entries ZS2, ZS3, ..., ZSN, ZSN+1. In this case, it is assumed that the repetition time TWN+1 is stored in a special column, though for greater clarity this is not shown in the figure. The address list of the new network access server ZSN+1 is still empty at this time point, however. Next, 10 second messages N21, N22, ..., N2N are sent from the network access servers of the group ZS1, ZS2, ..., ZSN to the distribution list address. However, each of these messages also contains an identification of the group GI and the repetition times TW1, TW2, ..., TWN that have been assigned to the network 15 access servers. As a next step, an address list is now also constructed in the new network access server ZSN+1. This list includes the entries ZS1, ZS2, ..., ZSN.

Figure 3 shows that each network access server ZS in the group

20 ZS1, ZS2, ..., ZSN+1 sends a second message N21, N22, ..., N2N+1 to the other network access servers ZS in the group in a periodically recurring manner, specifically at the repetition times TW1, TW2, ..., TWN+1. In this case, a network access server ZS is deleted from the address lists of the other 25 network access servers ZS in this group if the second message N2 is not received from said network access server before the expiry of the repetition time TW. This condition does not occur in the example shown, however. Therefore the address lists remain unchanged.

30

Figure 4 shows the logging off of the n-th network access server ZSN from the group of network access servers ZS. For this, a third message N3, which again includes the identification of the group GI and an assigned repetition time

TWN, is sent via the distribution list address to the other network access servers ZS in the group. Upon receipt of this message N3, the address list there is adapted accordingly, i.e. the n-th access server ZSN is deleted.